

# Asynchronous Multi-Context Systems<sup>1</sup>

Stefan Ellmauthaler and Jörg Pührer<sup>2</sup>

**Abstract.** In this work, we present *asynchronous multi-context systems* (aMCSs), which provide a framework for loosely coupling different knowledge representation formalisms that allows for online reasoning in a dynamic environment. Systems of this kind may interact with the outside world via input and output streams and may therefore react to a continuous flow of external information. In contrast to recent proposals, contexts in an aMCS communicate with each other in an asynchronous way which fits the needs of many application domains and is beneficial for scalability. The federal semantics of aMCSs renders our framework an integration approach rather than a knowledge representation formalism itself. We illustrate the introduced concepts by means of an example scenario dealing with rescue services. In addition, we compare aMCSs to reactive multi-context systems and describe how to simulate the latter with our novel approach.

## 1 Introduction

Research in the field of knowledge representation (KR) has originated a plethora of different languages and formats. Based on these formal concepts a wealth of tools has emerged (e.g., ontologies, triple-stores, modal logics, temporal logics, nonmonotonic logics, logic programs under nonmonotonic answer set semantics, ...). In a “connected world” it is desirable not to spread out information over different applications but to have it available for every application if need be. Expressing all the knowledge usually represented in specifically tailored languages in a universal language would be too hard to achieve from the point of view of complexity as well as the troubles arising from the translation of the representations. Instead, a framework seems desirable that integrates multiple existing formalisms in order to represent every piece of knowledge in the language that is most appropriate for it.

Another aspect that has received little attention in the development of many KR formalisms is that in a variety of applications, knowledge is provided in a constant flow of information and it is desired to reason over this knowledge in a continuous manner. Many formalisms are conceptually one-shot formalisms: given a knowledge base, the user triggers the computation of a result (e.g., the answer to a query). In this paper we aim at using KR formalisms in an *on-line* fashion as it has been done in recent works, e.g., on stream data processing and querying [11, 10], stream reasoning with answer set programming [7], and forgetting [9, 5].

To address the demand for an integration of heterogeneous knowledge representation formalisms together with the awareness of a continuous flow of knowledge over time, reactive multi-context sys-

tems (rMCSs) [4] and evolving multi-context systems (eMCSs) [8] where proposed. Both frameworks are based on the ideas of managed multi-context systems (mMCSs) [3] which combine multiple *contexts* which can be seen as representations of different formalisms. The semantics of rMCSs and eMCSs are based on the notion of an *equilibrium* which realises a tight semantical integration of the different context formalisms which is in many applications not necessary. Due to reasoning over all contexts, the whole computation is necessarily synchronous as the different contexts have to agree on common beliefs for establishing equilibria.

Many real world applications which utilise communication between different services use asynchronous communication protocols (e.g., web services) and compute as soon as they have appropriate information about the problem they have to address. Therefore, we introduce *asynchronous multi-context systems* (aMCSs), a framework for loosely coupled knowledge representation formalisms and services. It still provides the capabilities to express different knowledge representation languages and the translation of information from one formalism to another. In addition, aMCSs are also aware of continuous streams of information and provide ways to model the asynchronous exchange of information. To communicate with the environment, they utilise input and output streams.

We will illustrate aMCSs using the example of a task planner for medical rescue units. Here, we assume a scenario where persons are calling an emergency response team to report incidents and the employee needs to collect all relevant information about the case. Afterwards, the case needs to be classified and available resources (e.g., free ambulances, ...) have to be assigned to the emergencies. In addition, current traffic data as well as the estimated time of arrival should be considered by another employee, the dispatcher. Our proposed aMCS that we understand as a recommender-system for the emergency response employee as well as to the dispatcher, incorporates different contexts like a medical ontology, a database with the current state of the ambulances, or a navigation system which is connected to a traffic density reporter. We want to stress that this might be one application where it would be a great gain for the overall system by allowing asynchronous computation and communication such that it is not necessary to wait for all other contexts (e.g., it would be unnecessary to wait for the recommendation of a plan of action for the dispatcher during the treatment of an emergency call).

The remainder of this paper is structured as follows. At first we will give a short background on concepts we need. In Section 3, we extend the basic ideas of MCS to propose our new notion of aMCS for modelling asynchronous interaction between coupled knowledge representation formalisms and formally characterise its behaviour over time. The subsequent section presents an example scenario, where asynchronous computation and a reactive response to different events is needed. Section 5 compares aMCSs to rMCSs and shows how the latter can be simulated by the former. Section 6 concludes

<sup>1</sup> This work has been partially supported by the German Research Foundation (DFG) under grants BR-1817/7-1 and FOR 1513.

<sup>2</sup> Institute of Computer Science, Leipzig University, Germany, email: {ellmauthaler,puehrer}@informatik.uni-leipzig.de

this paper with a discussion including an outlook on future work.

## 2 Preliminaries

We base our approach on the underlying ideas of mMCSs [3] which extend heterogeneous multi-context systems (MCSs) [2] by a management layer. It allows for complex updates and revisions of knowledge bases and is realised by a management function that provides the updates for each equilibrium. Despite we build on mMCSs, they differ substantially in some aspects from the formalism we introduce in this work for reasons intrinsic to the asynchronous approach (cf. Section 5). Consequently, we only reuse basic notions from the original work and refer the interested reader to the paper of Brewka *et al.* [3] for full details on mMCS.

Like mMCS, aMCSs build on an abstract notion of a *logic suite* which can be seen as an abstraction of different formalisms for knowledge representation. A logic suite is a triple  $\mathcal{LS} = \langle \mathcal{KB}, \mathcal{BS}, \mathcal{ACC} \rangle$ , where  $\mathcal{KB}$  is the set of admissible knowledge bases (KBs) of  $\mathcal{LS}$ . Each knowledge base is a set of formulas that we do not further specify.  $\mathcal{BS}$  is the set of possible belief sets of  $\mathcal{LS}$ , whose elements are beliefs. A semantics for  $\mathcal{LS}$  is a function  $\mathcal{ACC} : \mathcal{KB} \rightarrow 2^{\mathcal{BS}}$  assigning to each KB a set of acceptable belief sets. Using a semantics with potentially more than one acceptable belief set allows for modelling non-determinism, where each belief set corresponds to an alternative solution. Finally,  $\mathcal{ACC}$  is a set of semantics for  $\mathcal{LS}$ . We denote  $\mathcal{KB}, \mathcal{BS}$ , respectively,  $\mathcal{ACC}$  by  $\mathcal{KB}_{\mathcal{LS}}, \mathcal{BS}_{\mathcal{LS}}$ , respectively,  $\mathcal{ACC}_{\mathcal{LS}}$ .

The motivation behind having multiple semantics for one formalism is that in our framework, the semantics of a formalism can be changed over time. While it is probably rarely the case that one wants to switch between different families of semantics during a run, e.g., from the stable-model semantics to the well-founded semantics of logic programs other switches of semantics are quite natural to many applications: we use different semantics to express different reasoning modes or to express different queries, i.e.,  $\mathcal{ACC}_1$  returns belief sets answering a query  $q_1$ , whereas  $\mathcal{ACC}_2$  answers query  $q_2$ ;  $\mathcal{ACC}_3$ , in turn, could represent the computation of all solutions to a problem, whereas at some point in time one could be interested in using  $\mathcal{ACC}_4$  that only computes a single solution. For instance one that is optimal with respect to some criterion.

## 3 Asynchronous Multi-Context Systems

An aMCS is built up by multiple contexts which are defined next and which are used for representing reasoning units. We assume a set  $\mathcal{N}$  of names that will serve as labels for sensors, contexts, and output streams.

**Definition 1** A context is a pair  $C = \langle n, \mathcal{LS} \rangle$  where  $n \in \mathcal{N}$  is the name of the context and  $\mathcal{LS}$  is a logic suite.

For a given context  $C = \langle n, \mathcal{LS} \rangle$  we denote  $n$  and  $\mathcal{LS}$  by  $n_C$  and  $\mathcal{LS}_C$ , respectively.

**Definition 2** An aMCS (of length  $n$  with  $m$  output streams) is a pair  $M = \langle C, O \rangle$ , where  $C = \langle C_1, \dots, C_n \rangle$  is an  $n$ -tuple of contexts and  $O = \langle o_1, \dots, o_m \rangle$  with  $o_j \in \mathcal{N}$  for each  $1 \leq j \leq m$  is a tuple containing the names of the output streams of  $M$ .

By  $\mathcal{N}(M)$  we denote the set  $\{n_{C_1}, \dots, n_{C_n}, o_1, \dots, o_m\}$  of names of contexts and output streams of  $M$ .

A context in an aMCS communicates with other contexts and the outside world by means of streams of data. In particular, we assume

that every context has an input stream on which information can be written from both external sources (we call them sensors) and internal sources (i.e., other contexts). For the data in the communication streams we assume a communication language  $\mathcal{IL}$  where every  $i \in \mathcal{IL}$  is an abstract *piece of information*. In our framework, the data in the input stream of a context and the data in output streams are modelled by information buffers that are defined in the following.

**Definition 3** A data package is a pair  $d = \langle s, I \rangle$ , where  $s \in \mathcal{N}$  is either a context name or a sensor name, stating the source of  $d$ , and  $I \subseteq \mathcal{IL}$  is a set of pieces of information. An information buffer is a sequence of data packages.

As we assume that data is asynchronously passed to a context on its input stream, it is natural that not all information required for a computation is available at all times. Consequently, we need means to decide whether a computation should take place, depending on the current KB and the data currently available on the stream, or whether the context has to wait for more data. In our framework, this decision is made by a computation controller as defined next.

**Definition 4** Let  $C = \langle n, \mathcal{LS} \rangle$  be a context. A computation controller for  $C$  is a relation  $cc$  between a KB  $\mathcal{KB} \in \mathcal{KB}_{\mathcal{LS}}$  and a finite information buffer.

Thus, if  $\langle \mathcal{KB}, \text{ib} \rangle \in cc$  then a computation should take place, whereas  $\langle \mathcal{KB}, \text{ib} \rangle \notin cc$  means that further information is required before the next computation is triggered in the respective context.

In contrast to the original definition of multi-context systems [1] and extensions thereof, we do not make use of so-called *bridge rules* as a means to communicate: a bridge rule defines which information a context should obtain based on the results of all the contexts of a multi-context system. In the asynchronous approach, we do not have (synchronised) results of all contexts available in general. As a consequence we use another type of rules, called *output rules*, that define which information should be sent to another context or an output stream, based on a result of a single context.

**Definition 5** Let  $C = \langle n, \mathcal{LS} \rangle$  be a context. An output rule  $r$  for  $C$  is an expression of the form

$$\langle n, i \rangle \leftarrow b_1, \dots, b_j, \text{not } b_{j+1}, \dots, \text{not } b_m, \quad (1)$$

such that  $n \in \mathcal{N}$  is the name of a context or an output stream,  $i \in \mathcal{IL}$  is a piece of information, and every  $b_\ell$  ( $1 \leq \ell \leq m$ ) is a belief for  $C$ , i.e.,  $b_\ell \in S$  for some  $S \in \mathcal{BS}_{\mathcal{LS}}$ .

We call  $n$  the *stakeholder* of  $r$ ,  $\langle n, i \rangle$  the head of  $r$  denoted by  $hd(r)$  and  $b_1, \dots, b_j, \text{not } b_{j+1}, \dots, \text{not } b_m$  the body  $bd(r)$  of  $r$ . Moreover, we say that  $r$  is active under  $S$ , denoted by  $S \models bd(r)$ , if  $\{b_1, \dots, b_j\} \subseteq S$  and  $\{b_{j+1}, \dots, b_m\} \cap S = \emptyset$ .

Intuitively, the stakeholder is a reference to the addressee of information  $i$ .

**Definition 6** Let  $C = \langle n, \mathcal{LS} \rangle$  be a context,  $OR$  a set of output rules for  $C$ ,  $S \in \mathcal{BS}_{\mathcal{LS}}$  a belief set, and  $n' \in \mathcal{N}$  a name. Then, the data package

$$d_C(S, OR, n') = \langle n, \{i \mid r \in OR, hd(r) = \langle n', i \rangle, S \models bd(r)\} \rangle$$

is the output of  $C$  with respect to  $OR$  under  $S$  relevant for  $n$ .

Compared to previous notions of multi-context systems, contexts in our setting only specify which formalisms they use but they do

not contain knowledge bases, the concrete semantics to use, and communication specifications. The reason is that for aMCSs these may change over time. Instead, we wrap concepts that are subject to change during runtime in the following notion of a *configuration*.

**Definition 7** Let  $C = \langle n, \mathcal{LS} \rangle$  be a context. A configuration of  $C$  is a tuple  $cf = \langle KB, ACC, ib, cm \rangle$ , where  $KB \in \mathcal{KB}_{\mathcal{LS}}$ ,  $ACC \in \mathcal{ACC}_{\mathcal{LS}}$ ,  $ib$  is a finite information buffer, and  $cm$  is a context management for  $C$  which is a triple  $cm = \langle cc, cu, OR \rangle$ , where

- $cc$  is a computation controller for  $C$ ,
- $OR$  is a set of output rules for  $C$ , and
- $cu$  is a context update function for  $C$  which is a function that maps an information buffer  $ib = d_1, \dots, d_m$  and an admissible knowledge base of  $\mathcal{LS}$  to a configuration  $cf' = \langle KB', ACC', ib', cm' \rangle$  of  $C$  with  $ib' = d_k, \dots, d_m$  for some  $k \geq 1$ .

We write  $cc_{cm}$ ,  $cu_{cm}$ , and  $OR_{cm}$  to refer to the components of a given context management  $cm = \langle cc, cu, OR \rangle$ . The context management is the counterpart of a *management function* of an rMCS, that computes an update of the knowledge base of a context given the results of bridge rules of the context.

In Section 2 we already discussed why we want to change semantics over time. Allowing also for changes of output rules can be motivated with applications where it should be dynamically decided where to direct the output of a context. For example, if a particular subproblem can be solved by two contexts  $C_1$  and  $C_2$  and it is known that some class of instances can be better solved by  $C_1$  and others by  $C_2$ . Then a third context that provides an instance can choose whether  $C_1$  or  $C_2$  should carry out the computation by adapting its output rules. Dynamically changing output rules and semantics could require adjustments of the other components of the context management. Thus, it makes sense that also computation controllers and context update functions are subject to change for the sake of flexibility.

**Definition 8** Let  $M = \langle \langle C_1, \dots, C_n \rangle, \langle o_1, \dots, o_m \rangle \rangle$  be an aMCS. A configuration of  $M$  is a pair

$$Cf = \langle \langle cf_1, \dots, cf_n \rangle, \langle ob_1, \dots, ob_m \rangle \rangle,$$

where

- for all  $1 \leq i \leq n$   $cf_i = \langle KB, ACC, ib, cm \rangle$  is a configuration for  $C_i$  and for every output rule  $r \in OR_{cm}$  we have  $n \in \mathcal{N}(M)$  for  $\langle n, i \rangle = hd(r)$ , and
- $ob_j = \dots, d_{l-1}, d_l$  is an information buffer with a final element  $d_l$  that corresponds to the data on the output stream named  $o_j$  for each  $1 \leq j \leq m$  such that for each  $h \leq l$  with  $d_h = \langle n, i \rangle$  we have  $n = n_{C_i}$  for some  $1 \leq i \leq n$ .


Figure 1 depicts an aMCS  $M$  with three contexts and a configuration for  $M$ .

We next characterise the dynamic behaviour of an aMCS. For easier notation we stick to a discrete notion of time represented by integers.

**Definition 9** Let  $M = \langle \langle C_1, \dots, C_n \rangle, \langle o_1, \dots, o_m \rangle \rangle$  be an aMCS. A run structure for  $M$  is a sequence

$$R = \dots, Cf^t, Cf^{t+1}, Cf^{t+2}, \dots,$$

where  $t \in \mathbb{Z}$  is a point in time, and every  $Cf^{t'}$  in  $R$  ( $t' \in \mathbb{Z}$ ) is a configuration of  $M$ .



**Figure 1.** An aMCS with three contexts, three sensors on the left side, and three output streams on the right side. A solid line represents a flow of information from a context to its stakeholder streams, whereas a dashed line indicates sensor data written to the input buffer of a context.

We will sometimes use  $cf_i^t$  to denote the configuration of a context  $i$  that appears at time  $t$  in a given run structure in the context of a given aMCS. Similarly,  $ob_j^t$  refers to the information buffer representing the data in the output stream named  $o_j$ . Moreover, we write  $KB_i^t$ ,  $ACC_i^t$ ,  $ib_i^t$ , and  $cm_i^t$  to refer to the components of  $cf_i^t = \langle KB, ACC, ib, cm \rangle$ . We say that context  $C_i$  is *waiting* at time  $t$  if  $\langle KB_i^t, ib_i^t \rangle \notin cc_{cm_i^t}$ .

**From run structure to run** In aMCSs we take into account that the computation of the semantics of a knowledge base needs time. Moreover, in a computation of our framework, different belief sets may become available at different times and verifying the non-existence of further belief sets can also take time after the final belief set has been computed. In order to model whether a context is busy with computing, we introduce a boolean variable  $busy_i^t$  for each configuration  $cf_i^t$  in a run structure. Hence, context  $C_i$  is *busy* at time  $t$  iff  $busy_i^t$  is true. While a context is busy, it does not read new information from its input stream until every belief set has been computed and it has concluded that no further belief set exists.

After the computation of a belief set, the output rules are applied in order to determine which data is passed on to stakeholder contexts or output streams. These are represented by *stakeholder buffers*: An information buffer  $b$  is the stakeholder buffer of  $C_i$  (for  $n$ ) at time  $t$  if

- $b = ib_{i'}^t$  for some  $1 \leq i' \leq n$  such that  $n = n_{C_{i'}}$  is stakeholder of some output rule in  $OR_{cm_{i'}^t}$  or
- $b = ob_{j'}^t$  for some  $1 \leq j' \leq m$  such that  $n = o_{j'}$  is stakeholder of some output rule in  $OR_{cm_i^t}$ .

In order to indicate that a computation has finished we assume a dedicated symbol  $EOC \in \mathcal{IL}$  that notifies a context's stakeholder buffers about the end of a computation.

Next, we formally characterise the behaviour of aMCSs followed by a summary of its intuition.

**Definition 10** Let  $M$  be an aMCS of length  $n$  with  $m$  output streams and  $R$  a run structure for  $M$ .  $R$  is a run for  $M$  if the following conditions hold for every  $1 \leq i \leq n$  and every  $1 \leq j \leq m$ :

- (i) if  $cf_i^t$  and  $cf_i^{t+1}$  are defined,  $C_i$  is neither busy nor waiting at time  $t$ , then
  - $C_i$  is busy at time  $t + 1$ ,

$$- cf_i^{t+1} = cu_{cm_i^t}(ib_i^t, KB_i^t)$$

We say that  $C_i$  started a computation for  $KB_i^{t+1}$  at time  $t + 1$ .

(ii) if  $C_i$  started a computation for  $KB$  at time  $t$  then

- we say that this computation ended at time  $t'$ , if  $t'$  is the earliest time point with  $t' \geq t$  such that  $\langle n_{C_i}, EOC \rangle$  is added to every stakeholder buffer  $b$  of  $C_i$  at  $t'$ ; the addition of  $d_{C_i}(S, OR_{cm_i^{t'}}, n)$  to  $b$  is called an end of computation notification.
- for all  $t' > t$  such that  $cf_i^{t'}$  is defined,  $C_i$  is busy at  $t'$  unless the computation ended at some time  $t''$  with  $t < t'' < t'$ .
- if the computation ended at time  $t'$  and  $cf_i^{t'+1}$  is defined then  $C_i$  is not busy at  $t' + 1$ .

(iii) if  $C_i$  started a computation for  $KB$  at time  $t$  that ended at time  $t'$  then for every belief set  $S \in ACC_i^t$  there is some time  $t''$  with  $t \leq t'' \leq t'$  such that

- $d_{C_i}(S, OR_{cm_i^{t''}}, n)$  is added to every stakeholder buffer  $b$  of  $C_i$  for  $n$  at  $t''$ .

We say that  $C_i$  computed  $S$  at time  $t''$ . The addition of  $d_{C_i}(S, OR_{cm_i^{t''}}, n)$  to  $b$  is called a belief set notification.

- (iv) if  $ob_j^t$  and  $ob_j^{t+1}$  are defined and  $ob_j^t = \dots, d_{l-1}, d_l$  then  $ob_j^{t+1} = \dots, d_{l-1}, d_l, \dots, d_{l'}$  for some  $l' \geq l$ . Moreover, every data package  $d_{l''}$  with  $l < l'' \leq l'$  that was added at time  $t + 1$  results from an end of computation notification or a belief set notification.
- (v) if  $cf_i^t$  and  $cf_i^{t+1}$  are defined,  $C_i$  is busy or waiting at time  $t$ , and  $ib_i^t = d_1, \dots, d_l$  then we have  $ib_i^{t+1} = d_1, \dots, d_l, \dots, d_{l'}$  for some  $l' \geq l$ . Moreover, every data package  $d_{l''}$  with  $l < l'' \leq l'$  that was added at time  $t + 1$  either results from an end of computation notification or a belief set notification or  $n \notin \mathcal{N}(M)$  (i.e.,  $n$  is a sensor name) for  $d_{l''} = \langle n, i \rangle$ .

Condition (i) describes the transition from an idle phase to an ongoing computation. The end of such a computation is marked by an end of computation notification as introduced in Item (ii). Condition (iii) states that between the start and the end of a computation all belief sets are computed and stakeholders are notified. Items (iv) and (v) express how data is added to an output stream or to an input stream, respectively. Note that here, sensors and the flow of information from sensors to the input buffers of contexts are implicit. That is, data packages from a sensor may appear at the end of input buffers at all times and the only reference to a particular sensor is its name appearing in a data package.

Summarising the behaviour characterised by a run, whenever a context  $C$  is not busy, its context controller  $cc$  checks whether a new computation should take place, based on the knowledge base and the current input buffer of  $C$ . If yes, the current configuration of the context is replaced by a new one, computed by the context update function  $cu$  of  $C$ . Here, the new input buffer has to be a suffix of the old one and a new computation for the updated knowledge base starts. After an undefined period of time, belief sets are computed and based on the application of output rules of  $C$ , data packages are sent to stakeholder buffers. At some point in time, when all belief sets have been computed, an end of computation notification is sent to stakeholders, and the context is not busy anymore.

## 4 Scenario: Computer-Aided Emergency Team Management

Now we want to consider a scenario, where aMCSs may be used to describe the asynchronous information-exchange between different specialised reasoning systems. Our example deals with a recommender-system for the coordination and handling of ambulance assignments. The suggested aMCS supports decisions in various stages of an emergency case. It gives assistance during the rescue call, helps in assigning priorities and rescue units to a case, and assists in the necessary communication among all involved parties. The suggestions given by the system are based on different specialised systems which react to sensor readings. Moreover, the system can tolerate and incorporate overriding solutions proposed by the user that it considers non-optimal.

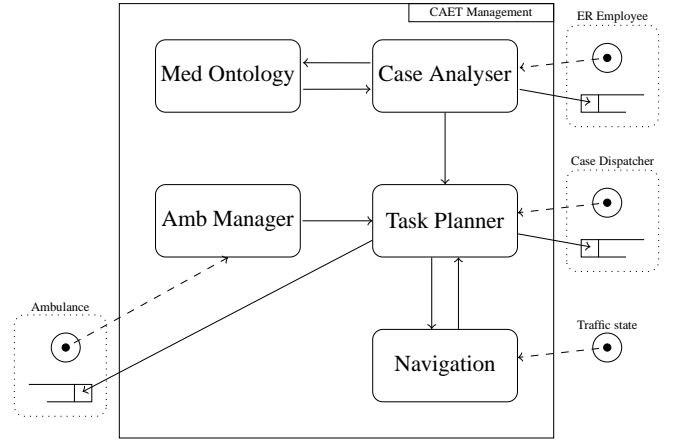


Figure 2. The Computer-Aided Emergency Team Management aMCS

Figure 2 depicts the example aMCS which models such a Computer-Aided Emergency Team Management System (CAET Management System). Note that interaction with a human (e.g., EM employee) is modelled as a pair containing an input stream and an output stream. The system consists of the following contexts:

**Case Analyser (ca)** This context implements a computer-aided call handling system which assists an emergency response employee (ER employee) during answering an emergency call. The system utilises reasoning methods to choose which questions need to be asked based on previous answers. In addition, it may check whether answers are inconsistent (e.g., amniotic sac bursts when the gender is male). For these purposes the case analyser context may also consult a medical ontology represented by another context. The communication with the ER employee is represented, on the one hand, as a sensor that reads the input of the employee and, on the other hand, by an output stream which prints the questions and results on a computer screen.

During the collection of all the important facts for this emergency case, the analyser computes the priority of the case and passes it to the task planner.

**Med Ontology (mo)** This medical ontology can be realised, e.g., by a description logic reasoner which handles requests from the case analyser and returns more specific knowledge about ongoing cases. This information may be used for the prioritisation of the importance of a case.

**Task Planner (tp)** This context keeps track of emergency cases. Based on the priority and age of a case and the availability and position of ambulances it suggests an efficient plan of action for the ambulances to the (human) case dispatcher (cd). The dispatcher may approve some of the suggestions or all of them. If the dispatcher has no faith in the given plan of action, she can also alter it at will. These decisions are reported back to the planning system such that it can react to the alterations and provide further suggestions. Based on the final plan, the task planner informs the ambulance about their new mission.

The knowledge base of the context is an answer-set program for reasoning about a suggested plan. It gets the availability and position of the ambulances by the ambulance manager. In addition, the cases with their priority are provided by the case analyser. With this information, the task planner gives the locations of the ambulances together with the target locations of the cases to a navigation system which provides the distances (i.e., the estimated time of arrival (ETA)) of all the ambulances to all the locations.

**Amb Manager (am)** The ambulance manager is a database, which keeps track of the status and location of ambulance units. Each ambulance team reports its status (e.g., to be on duty, waiting for new mission, ...) to the database (modelled by the sensor “Ambulance” (amb)). Additionally, the car periodically sends GPS-coordinates to the database. These updates will be pushed to the task planner.

**Navigation (na)** This part of the aMCS gets traffic information (e.g., congestions, roadblocks, construction zones, ...) to predict the travel time for each route as accurate as possible. The task planner may push a query to the navigation system, which consists of a list of locations of ambulance units and a list of locations of target areas. Based on all the given information this context will return a ranking for each target area, representing the ETAs for each ambulance.

Now we want to have a closer look on the instantiation details of some aspects of our example. At first we investigate the *cc* relation of the case analyser. It allows for the computation of new belief sets whenever the ER employee pushes new information to the analyser. In addition, it will also approve of a new computation if the medical ontology supplies some requested information. Recall that the case analyser also assigns a priority to each case and that we want to allow the employee to set the priority manually. Let us suppose that such a manual override occurs and that the case analyser has an ongoing query to the medical ontology. Due to the manual priority assignment, the requested information from the ontology is no longer needed. Therefore, it would be desirable that *cc* does not allow for a recomputation if all conclusions of the ontology are only related to the manually prioritised case. With the same argumentation in mind, the context update function *cu* will also ignore this information on the input stream. This kind of behaviour may need knowledge about past queries which can be provided by an additional output rule for the case analyser which feeds the relevant information back to the context.

Next, we will have a look at the task planner that is based on answer-set programming. We will only present parts of the program, to show how the mechanics are intended to work. To represent the incoming information on the input stream, the following predicates can be used:

**case(caseid,loc,priority)** represents an active case (with its location and priority) which needs to be assigned to an ambulance.

**avail(amb,loc)** states the location of an available ambulance.

**eta(caseid,amb,value)** provides the estimated time of arrival for a unit at the location of the target area of the case.

**assign(amb,caseid)** represents the assignment of an ambulance to a case by the dispatcher.

These predicates will be added by the context update function to the knowledge base if corresponding information is put on the input stream of the context. Based on this knowledge, the other components of the answer-set program will compute the belief sets (e.g., via the stable model semantics). Note that an already assigned ambulance or case will not be handled as an available ambulance or an active case, respectively. In addition, *cu* can (and should) also manage forgetting of no longer needed knowledge. For our scenario it may be suitable to remove all *eta*, *avail* and *case* predicates when the cases or the unit is assigned. The *assign* predicate can be removed when the ambulance manager reports that the assigned ambulance is available again.

The set OR of output rules of the task planner could contain the following rules:<sup>3</sup>

$$\begin{aligned} \langle \text{cd}, \text{assign}(A, C) \rangle &\leftarrow \text{sugassignment}(A, C) \\ \langle \text{na}, \text{queryA}(L) \rangle &\leftarrow \text{avail}(A), \text{not assign}(A, \_), \text{loc}(A, L) \\ \langle \text{na}, \text{queryC}(L) \rangle &\leftarrow \text{case}(C, P), \text{loc}(A, L), \text{not assign}(A, \_) \\ \langle \text{amb}, \text{assigned}(A, C) \rangle &\leftarrow \text{assign}(A, C) \end{aligned}$$

The first rule informs the case dispatcher (cd) about a suggested assignment that has been computed by the answer-set program. Rules two and three prepare lists of ambulances and cases for querying the navigation context. Recall that the latter needs a list of ambulance locations (generated by rule two) and a list of target area locations (generated by rule three). Also keep in mind that for each belief set a data package with all information for one context or output stream is constructed. So the whole list of current target areas and free ambulance units will be passed to the navigation context at once. The last rule notifies the ambulance team that it has been assigned to a specific case.

Related to this example we want to mention privacy aspects as a real world policy which is especially important to applications in public services and health care. As the multi-context system is a heterogeneous system with different contexts, a completely free exchange of data may be against privacy policies. This issue can be addressed by the adequate design of output rules, which can also be altered with respect to additional information in the input stream (e.g., some context gains the permission to receive real names instead of anonymous data). So each context may decide by its own which parts of the belief sets are shared and exchanged with other contexts.

Another interesting aspect about aMCSs is the possibility to easily join two aMCSs together, outsource a subset of contexts in a new aMCS, or to view an aMCS as an abstract context for another aMCS in a modular way. This can be achieved due to the abstract communication by means of streams. With respect to our scenario there could be some aMCS which does the management of resources for hospitals (e.g., free beds with their capabilities). The task planner might communicate with this system to take the needed services for a case into account (e.g., intensive care unit) and informs the hospital via these streams about incoming patients. It would be easy to join both aMCSs together to one big system or to outsource some contexts as input sensors paired with an output stream. In addition, one may also combine different contexts or a whole aMCS to one abstract context

<sup>3</sup> Keep in mind that in an actual implementation one may want to provide further information via communication.

to provide a dynamic granularity of information about the system and to group different reasoning tasks together.

## 5 Relation to Reactive Multi-Context Systems

In this section we want to address differences and communalities between aMCSs and rMCSs [4] as both are types of multi-context systems that work in an online fashion and can react to external information. Runs of rMCSs are based on equilibria which are collections of belief sets—one for each context—on which, intuitively, all of the contexts have to agree. Thus, equilibria can be seen as a tight integration approach in which the semantics of the individual contexts are interdependent. However, the high level of integration also comes at the price that the different contexts must wait for each other for the computation of each equilibrium, i.e., they are synchronised. In aMCSs, on the other hand, the coupling of the semantics is much looser—communication between contexts only works via data packages that are sent to another context after a computation and not via a higher-level common semantics for multiple contexts. But as a benefit, each context can run at its own pace which is useful in settings where there is a context that requires much more time for evaluating its semantics than others.

A further difference is the role of non-determinism in the semantics of aMCSs and rMCSs. An equilibrium in an aMCS consists of a single belief set for each context. Hence, as aMCSs also use a multiple belief set semantics, there may also be multiple equilibria as a source of non-determinism at each step in a run. For aMCSs, all belief sets of a context are computed in a consecutive way (we assume that if only a single belief set is desired than the semantics of the respective context should be adapted accordingly by the knowledge engineer). Nevertheless, there is also a source of non-determinism in the case of aMCSs caused by the undefined duration of computations.

Regarding the computational complexity of the two frameworks, the computation of an equilibrium requires to guess an equilibrium candidate first before the semantics of the context is computed which is expensive regarding runtime when put to practice. In theory, this guess does not add extra complexity if the context semantics is already NP-hard (as shown in [4]) because it can be combined with the guesses required in the contexts. However, this trick cannot be used in implementations that uses black boxes for computing context semantics. On the other hand, aMCSs do not add substantial computational requirements to the effort needed for computing context semantics. In particular, aMCSs are scalable as adding a further context has no direct influence on how the semantics of the other contexts are computed but can only influence the input they get.

Both, aMCSs and rMCSs are very general frameworks that allow for simulating Turing machines and thus for performing multi-purpose computations even if only very simple context formalisms are used (if the length of a run is not restricted). In this sense the approaches are equally expressive. Moreover, when allowing for arbitrary contexts one could trivially simulate the other by including it as a context. Despite the existence of these straightforward translations, we next sketch how we simulate an rMCS with an aMCS using a more direct translation, as this gives further insight into the differences of the two frameworks. Moreover, it demonstrates a way to implement rMCSs by means of aMCSs. For every context  $C_i$  of a given rMCS  $M_r$ , we introduce three contexts in the aMCS  $M_a$  that simulates  $M_r$ :

- a context  $C_i^{kb}$  that stores the current knowledge base of the context,

- a context  $C_i^{kb'}$  in which a candidate for an updated knowledge base can be written and its semantics can be computed, and
- a management context  $C_i^m$  that implements the bridge rules, and the management function of the context.

There are three further contexts:

- $C^{obs}$  receives sensor data and distributes it to every context  $C_i^m$  where  $C_i$  depends on the respective sensor. The context is also responsible for synchronisation: for each sensor, new sensor data is only passed on after an equilibrium has been computed.
- $C^{guess}$  guesses equilibrium candidates for  $M$  and passes them to the management contexts  $C_i^m$ . Based on that and the information from  $C^{obs}$ ,  $C_i^m$  computes an update  $kb'_i$  of the knowledge base in  $C_i^{kb}$  and stores  $kb'_i$  in  $C_i^{kb'}$ . The latter context then computes the semantics of  $kb'_i$  and passes it to the final context
- $C^{check}$  that compares every belief set it receives with the equilibrium candidate (that it also receives from  $C^{guess}$ ). If a matching belief set has been found for each context of  $M_r$ , the candidate is an actual equilibrium. In this case  $C^{check}$  sends the equilibrium to an output stream and notifies the other contexts about the success.

In case of a success, every context  $C_i^m$  replaces the knowledge base in  $C_i^{kb}$  by  $kb_i$  and a next iteration begins. In case no equilibrium was found but one of the  $C_i^{kb'}$  contexts has finished its computation,  $C^{check}$  orders  $C^{guess}$  to guess another equilibrium candidate.

## 6 Related Work and Discussion

A concept similar to output-rules has been presented in the form of reactive bridge rules [6]. There the flow of information is represented by rules which add knowledge to the input streams of other contexts. Which information is communicated to other contexts is also determined by the local belief set of each context.

Note that evolving multi-context systems [8] follow a quite similar approach as rMCSs and hence the relation of aMCSs to rMCSs sketched in the previous section also applies in essence to this approach.

The system `clingo` [7] is a reactive answer-set programming solver. It utilises TCP/IP ports for incoming input streams and does also report the resulting answer sets via such a port. It provides means to compute different semantics and can keep learned structures and knowledge from previous solving steps. Although there are no output rules or input stream pre-processing as in aMCSs, the system features embedded imperative programming languages which may be helpful to model some of the presented concepts of this paper.

In general, the tasks performed by a context management can be realised by different formalisms (e.g., imperative scripting languages or declarative programming). Here, it seems likely that different languages can be the most appropriate management language, depending on the type of context formalism and the concrete problem domain. A feature that is not modelled in our proposal but that is potentially useful and we intend to consider in the future is to allow for aborting computations. Moreover, we want to study modelling patterns and best practices for aMCSs design for typical application settings and compare different inter-context topologies and communication strategies.

The next natural step towards an implementation is an analysis of how existing tools such as `clingo` could be used for a realisation. It is clear that such formalisms can be used as a context formalism. Moreover, we are interested in how reactive features of

`clingo` (e.g., iterative computation, on-demand grounding, online-queries, ...) relate to aMCS concepts (e.g., `cc`, `ib`, ...) and whether the system can be described in terms of an aMCS.

## REFERENCES

- [1] Gerhard Brewka and Thomas Eiter, ‘Equilibria in heterogeneous non-monotonic multi-context systems’, in *AAAI’07*, pp. 385–390, (2007).
- [2] Gerhard Brewka, Thomas Eiter, and Michael Fink, ‘Nonmonotonic multi-context systems: A flexible approach for integrating heterogeneous knowledge sources’, in *Logic Programming, Knowledge Representation, and Nonmonotonic Reasoning*, 233–258, Springer, (2011).
- [3] Gerhard Brewka, Thomas Eiter, Michael Fink, and Antonius Weinzierl, ‘Managed multi-context systems’, in *IJCAI’11*, pp. 786–791, (2011).
- [4] Gerhard Brewka, Stefan Ellmauthaler, and Jörg Pührer, ‘Multi-context systems for reactive reasoning in dynamic environments’, in *Proc. ECAI’14*, (2014). To appear.
- [5] Fu-Leung Cheng, Thomas Eiter, Nathan Robinson, Abdul Sattar, and Kewen Wang, ‘Lpforget: A system of forgetting in answer set programming’, in *Proc. AUSAI’06*, eds., Abdul Sattar and Byeong Ho Kang, volume 4304 of *LNCS*, pp. 1101–1105. Springer, (2006).
- [6] Stefan Ellmauthaler, ‘Generalizing multi-context systems for reactive stream reasoning applications’, in *Proc. ICCSW’13*, pp. 17–24, (2013).
- [7] Martin Gebser, Torsten Grote, Roland Kaminski, Philipp Obermeier, Orkunt Sabuncu, and Torsten Schaub, ‘Stream reasoning with answer set programming: Preliminary report’, in *Proc. KR’12*, (2012).
- [8] Ricardo Gonçalves, Matthias Knorr, and João Leite, ‘Evolving multi-context systems’, in *Proc. ECAI’14*, (2014). To appear.
- [9] Jérôme Lang and Pierre Marquis, ‘Reasoning under inconsistency: A forgetting-based approach’, *Artif. Intell.*, **174**(12-13), 799–823, (2010).
- [10] Danh Le-Phuoc, Josiane Xavier Parreira, and Manfred Hauswirth, ‘Linked stream data processing’, in *Proc. RW’12*, eds., Thomas Eiter and Thomas Krennwallner, volume 7487 of *LNCS*, pp. 245–289. Springer, (2012).
- [11] Carlo Zaniolo, ‘Logical foundations of continuous query languages for data streams’, in *Proc. Datalog 2.0*, pp. 177–189, (2012).